

REMARKS

Claims 1-24 are pending in the present application. No amendments to the claims are made by this Response. Reconsideration of the claims is respectfully requested.

I. Telephone Interview

Applicants thank Examiner Bates for the courtesies extended to Applicants' representative during the March 29, 2004 telephone interview. During the interview, Applicants' representative discussed the many distinctions of the present claims over the Lee reference. Applicants' representative submitted arguments that the Examiner cannot simply reject the general concept of suspending "something," performing a modification, and then placing the suspended "something" into operation again but must instead examine the actual features recited in the claims. Applicants' representative further argued that none of the specific features of the dependent claims are taught or suggested anywhere in Lee. Examiner Bates agreed that the Lee reference does not teach a send queue, and thus cannot teach suspending a send queue or placing the send queue back into an operational state following application of a modification to a network. However, Examiner Bates requested that detailed arguments be provided in support of the differences expressed during the telephone interview. Thus, the substance of the interview is summarized in the following remarks.

II. 35 U.S.C. § 102, Alleged Anticipation

The Office Action rejects claims 1, 3-9, 12-17, and 20-24 under 35 U.S.C. § 102(b) as being anticipated by Lee (U.S. Patent No. 4,951,225). This rejection is respectfully traversed.

As to independent claims 1, 9 and 17, the Office Action states:

Regarding claims 1, 9, and 17, Lee discloses a method for modifying a network without tearing down existing connections (Column 3, lines 64-68; Column 5, lines 50-53), comprising: placing a send queue that is to be affected by a modification to the network into a suspended

state (Column 4, lines 5-8; Column 6, lines 32-35); applying the modification to the network (Column 4, lines 13-17; lines 20-21); and placing the send queue back into an operational state after applying the modification to the network (Column 12, lines 15-17; Column 13, lines 16-19).

Office Action dated December 31, 2003, page 2.

Claim 1, which is representative of the other rejected independent claims 9 and 17 with regard to similarly recited subject matter, reads as follows:

1. A method for modifying a network without tearing down existing connections, comprising:
 - placing a send queue that is to be affected by a modification to the network into a suspended state;
 - applying the modification to the network; and
 - placing the send queue back into an operational state after applying the modification to the network. (emphasis added)

A prior art reference anticipates the claimed invention under 35 U.S.C. § 102 only if every element of a claimed invention is identically shown in that single reference, arranged as they are in the claims. *In re Bond*, 910 F.2d 831, 832, 15 U.S.P.Q.2d 1566, 1567 (Fed. Cir. 1990). All limitations of the claimed invention must be considered when determining patentability. *In re Lowry*, 32 F.3d 1579, 1582, 32 U.S.P.Q.2d 1031, 1034 (Fed. Cir. 1994). Anticipation focuses on whether a claim reads on the product or process a prior art reference discloses, not on what the reference broadly teaches. *Kalman v. Kimberly-Clark Corp.*, 713 F.2d 760, 218 U.S.P.Q. 781 (Fed. Cir. 1983). Applicants respectfully submit that the Lee reference does not identically show every element of the claimed invention arranged as they are in the claims. Specifically, Lee does not teach send queues and thus, cannot teach placing a send queue in a suspended state or placing the send queue back into an operational state, as recited in claims 1, 9 and 17.

Lee is directed to a updating working memories of nodes in a Rete network. A Rete network is an artificial intelligence inference engine that permits matching a large collection of patterns to a large collection of objects is one operation of an inference engine. An inference engine is the reasoning component in artificial intelligence and expert systems. Inference engines typically contain three components or methods: (1)

matching a collection of objects to a given rule, (2) selecting a rule from a list of rules whose conditions are completely matched by objects, and (3) firing of the selected rule. The most time consuming phase in an inference engine is the matching process. For example, consider a system with a 1000 objects and 1000 rules, where each rule has three conditions. In a very simple implementation scheme, each rule is compared against all objects. This leads to over a trillion match operations for each execution cycle.

The Rete pattern matching algorithm is a complex algorithm that can achieve the result of this example with only a small fraction of the work (that is, much fewer than a trillion match operations). The Rete pattern matching algorithm implements a "divide and conquer" strategy to the matching component or method. The conditions that make up a given rule are tested. Failure of one condition prevents testing of subsequent Conditions. To perform the matching the Rete pattern matching algorithm uses an augmented discrimination network, referred to as a Rete network, compiled from the premise of the rule. Each object is first tested against conditions that involve only that object. Objects which pass these tests are then joined with other objects to perform tests requiring multiple objects. Collections of objects which pass all tests become candidates for selection and firing (see the Background of U.S. Patent No. 5, 276,776).

The Lee system and method improves upon the Rete network by providing a mechanism for updating the nodes within the Rete network in response to a change in a working memory element (WME). In other words, Lee is directed to solving the problem of correctly updating a Rete network, i.e. an inference engine (see column 7, lines 51-54). Basically, the Lee system provides a mechanism for stopping and resuming updates to working memories of nodes, i.e. skipping over nodes, based on processing of a change token in an alpha portion of the Rete network. The Lee system is used to update the rules that are applied by the nodes of the Rete network but to limit the propagation of the update to only those nodes that need to be updated. The nodes that do not receive the updates to the working memory, i.e. have a stop node that stops the propagation of the change token, are term "suspended" nodes (see column 4, lines 8-12).

The Rete network has nothing to do with send queues in a network. The Rete "network" is merely a data structure representation of a nodal network in which each node applies certain tests, or a portion of a rule, to work memory entries. The Rete

network does not have send queues and thus, cannot suspend a send queue or place a send queue back into an operational state.

Furthermore, nowhere in Lee is the phrase "send queue" ever mentioned. This is because, as stated above, the Rete network has nothing to do with send queues. It is apparent from the Office Action that the nodes of the Rete network are being equated to a send queue. However, the nodes of the Rete network do not perform any functions similar to that of a send queue. To the contrary, the nodes in the Rete network are merely data structures that represent portions of a rule to be applied to work memory entries (see column 5, lines 9-18).

While Lee may use the words "update", "suspend" and "resume" with regard to Rete network nodes, these operations are not performed with regard to a send queue. To the contrary, the "update" is to update rules applied by nodes in a Rete network, which merely apply rules to work memory entries. The "suspended nodes" are nodes that do not receive the updates because they are associated with a stop node. The "resume node" is merely a node that permits resuming propagation of a change token through the Rete network following a stop node. The combination of the stop node and the resume node permit portions of the Rete network to be skipped during the update, i.e. the propagation of the change token. The "update", "suspend" and "resume" terms, as they are used in Lee, have nothing to do with placing a send queue that is to be affected by a modification to a network into a suspended state or placing a send queue back into an operational state after applying a modification to a network.

Furthermore, even if the nodes of a Rete network were somehow considered equivalent to the send queue of the presently claimed invention, which it is not, Lee still does not teach placing a send queue that is to be affected by a modification to the network into a suspended state. To the contrary, Lee teaches that the "suspended" nodes are those nodes that are not updated. See column 4, lines 8-12 where it states that "all nodes which are not updated are members of a second portion of the network termed 'suspended' portion signifying that such predetermined nodes have not been updated to the latest work update." Thus, Lee specifically teaches to suspend nodes that are not to be affected by updates to the Rete network. This is the exact opposite of what is claimed.

that is if one were to somehow equate the nodes of a Rete network to a send queue, which would be completely against the actual teachings of the Lee reference.

The actual portions of the Lee reference cited by the Office Action as allegedly teaching the features of the present claims actually provide further evidence of the fact that Lee does not teach anything in the claims. For example, the Office Action alleges that placing a send queue that is to be affected by a modification to a network in a suspended state is taught at column 4, lines 5-8 and column 6, lines 32-35, which read as follows:

Stop nodes are designated in the network to stop the updating from reaching predetermined ones of the nodes; such stop nodes are updated. All nodes which are not updated are members of a second portion of the network termed "suspended" portion signifying that such predetermined nodes have not been updated to the latest network update.
(column 4, lines 5-12)

That is, stop nodes are introduced to enable a clear description of a partially updated RETE network. Such stop nodes may be simultaneously as a stop node preventing token propagation to predetermined successors while other successors are not subjected to the stopped updating.
(column 6, lines 30-35)

Nowhere in either of these sections is a send queue even mentioned, let alone placing a send queue that is to be affected by a modification to a network into a suspended state. To the contrary, the first cited section supports Applicants' assertion that the stop nodes in Lee prevent propagation of changes to certain nodes and that these nodes are called "suspended" nodes. The second cited section of Lee supports Applicants' assertion that the stop nodes prevent token propagation to predetermined successors, i.e. the suspended nodes. Thus, the "suspended" nodes in Lee are nodes that are not affected by the update to the Rete network, as opposed to the currently claimed invention where the send queue is suspended when it is a send queue that is to be affected by a modification to the network.

As a further example, the Office Action alleges that the feature of placing a send queue back into an operational state after applying a modification to a network is taught by Lee at column 12, lines 15-17 and column 13, lines 16-19, which read as follows:

Shadowing means that the associate result memory for the stop node is put in the state it would be before the stop node processed the suspended token
(column 12, lines 15-17)

This completes the updating operation for a generalized RETE with stop nodes. Program 11 now resumes program execution as provided in the prior art systems.
(column 13, lines 16-19).

Column 12, lines 15-17 merely states that the term "shadowing" is defined as the associate result memory for the stop node being put in a state that it would be before the stop node processed the suspended token. There is nothing in this statement that teaches to place a send queue back into an operational state after applying a modification to a network. There simply is no correlation between the cited section of the reference and anything in the claimed invention.

Furthermore, column 13, lines 16-19 does not teach anything regarding a send queue, placing a send queue in an operational state, or placing a send queue in an operational state after applying a modification to a network. Once again, there is simply no correlation between the cited section of the reference and anything in the claimed invention. These seem to be randomly selected and irrelevant portions of the reference that in fact have nothing to do with the claimed invention.

Thus, in view of the above, Applicants respectfully submit that Lee does not teach each and every feature of claims 1, 9 and 17 as is required under 35 U.S.C. § 102(b). At least by virtue of their dependency on claims 1, 9 and 17, respectively, Lee does not teach each and every feature of dependent claims 3-8, 12-16, and 20-24. Accordingly, Applicants respectfully request withdrawal of the rejection of claims 1, 3-9, 12-17, and 20-24 under 35 U.S.C. § 102(b).

Furthermore, Lee does not teach, suggest, or give any incentive to make the needed changes to reach the presently claimed invention. Lee actually has nothing to do with send queues, placing send queues in a suspended state, or placing send queues in an operational state after applying modifications to a network. Absent the Examiner pointing out some teaching or incentive to implement Lee such that Lee operates on send queues by placing send queues in a suspended state and placing send queues in an

operational state after applying modifications to a network, one of ordinary skill in the art would not be led to modify Lee to reach the present invention when the reference is examined as a whole. Absent some teaching, suggestion, or incentive to modify Lee in this manner, the presently claimed invention can be reached only through an improper use of hindsight using the Applicants' disclosure as a template to make the necessary changes to reach the claimed invention.

In addition to the above, Lee does not teach the specific features recited in the dependent claims 3-8, 12-16, and 20-24. For example, with regard to claim 3, Lee does not teach that while in a suspended state, write requests can be posted to a queue pair of the send queue, incoming message to the queue pair of the send queue are processed normally, and work requests submitted to the send queue are queued and are not processed. The Office Action fails to even address these features in the context of the 102(b) rejection but instead addresses these features in the rejection under 35 U.S.C. § 103(a), as discussed hereafter. Thus, the Office Action has failed to establish a case of anticipation with regard to the features of claim 3.

With regard to claims 4, 12 and 20, Lee does not teach that placing a send queue into a suspended state includes stopping processing of message in the send queue at a message boundary. As stated above, Lee has nothing to do with send queues. Moreover, nothing in Lee teaches or even suggests stopping processing of messages in a send queue at a message boundary. Lee does not even mention messages in a send queue or message boundaries, let alone stopping processing of message in a send queue at a message boundary.

The Office Action alleges that this feature is taught by Lee at column 6, lines 45-48 which reads as follows:

The tokens stopped at the stop nodes are therefore called "suspended tokens" to signify that they have been temporarily suspended from pushing, but need more pushing into the suspended RETE whenever the RETE is to be updated in a manner to use such predetermined successor nodes.

Nowhere in this or any other section of Lee is there any mention of messages in a send queue, message boundaries, or stopping processing of message in a send queue at a

message boundary. Just as with every other section of Lee cited by the Office Action, this portion of Lee is completely irrelevant to the actual claim features. Thus, the Office Action has failed to establish a case of anticipation with regard to claims 4, 12 and 20.

With regard to claims 5, 13 and 21, Lee does not teach sending a notification to a subnet manager that the send queue has been placed in a suspended state, wherein the modification to the network is applied in response to sending the notification. Again the Office Action points to a seemingly irrelevant portion of the Lee reference, i.e. column 12, line 57 to column 13, line 5, which reads as follows:

In the machine executed step 68, the shadowing is removed; i.e., the update implied by the shadow token is done to the associated result memory of the beta stop node. This may add, remove, or modify a token in the result memory. The illustrated machine executed step 69 pushes the just restored token in an ascending scan of the resume nodes in the suspended portion of the pattern matching network. This step is the second updating step of the two complementary updating steps. Here, ascending order is determined as described above, in terms of upward or downward reachability among disjoint subsets of mutually reachable resume nodes. Upon completion of step 69, the suspended token becomes finalized, permanently stored in the suspended portion of the RETE. The above three steps 67, 68 and 69 are repeated for all suspended tokens in all beta stop nodes in any order.

Where in here is there any mention of a subnet manager, a send queue, sending a notification to a subnet manager that a send queue has been placed in a suspended state, or that the modification to a network is applied in response to sending the notification to the subnet manager? There simply is no correlation between the cited section of the reference and the presently claimed invention other than the use of the word "suspended." Just because the reference uses the word "suspended" does not mean that the specific features of claim 5, 13 and 21 are taught by the reference. In actuality, as stated many times above, there is nothing in Lee that teaches anything remotely resembling the claimed features. Other than merely using some similar words such as "suspended", the reference is completely irrelevant.

The other dependent claims 6-8, 14-16, and 22-24 recite other features of the present invention that simply are not taught by the Lee reference. The Lee reference is

equality irrelevant with regard to the features of claims 6-8, 14-16 and 22-24 as it is with the claims discussed above. Thus, the Lee reference does not teach any of the features of the dependent claims 3-8, 12-16, and 20-24 and the rejection of these claims under 35 U.S.C. § 102(b) should be withdrawn.

III. 35 U.S.C. § 103, Alleged Obviousness

The Office Action rejects claims 2-3, 10-11, and 18-19 under 35 U.S.C. § 103 as being allegedly unpatentable over Lee. This rejection is respectfully traversed for at least the same reasons as set forth above with regard to independent claims 1, 9 and 17. Lee does not teach or suggest a send queue and thus, cannot teach or suggest any of the features of claims 1, 9 and 17.

In addition, as to the specific features of claims 2-3, 10-11, and 18-19, the Office Action states that simply because Lee allegedly teaches that "some conditions in the network, such as modifications to network nodes may result in invalid and unexpected results if the network processes tokens or packets while needing to be updating, or in the middle of updating" that somehow all of the specific features of claims 2-3, 10-11 and 18-19 are rendered obvious in some way. Claims 2, 10 and 18 specifically recite that the suspended state is a send queue drain state. As noted above, Lee does not even teach a send queue and thus, cannot teach a send queue drain state as recited in these claims. There is simply nothing in Lee that is equivalent to a send queue drain state despite the allegations made by the Office Action.

Claims 3, 11 and 19 recite that while in the suspended state, write requests can be posted to a queue pair of the send queue, incoming message to the queue pair of the send queue are processed normally, and work requests submitted to the send queue are queued and are not processed. Nowhere in Lee is there any teaching of a send queue, a queue pair, a send queue being part of a queue pair, messaging incoming to a queue pair, posting of requests to a queue pair, work requests being submitted to a send queue being queued but not processed, or any other requirement of claims 3, 11 and 19. Lee is completely irrelevant to the present claims and thus, does not teach these features. Furthermore, the Office Action has failed to specifically show where any of these

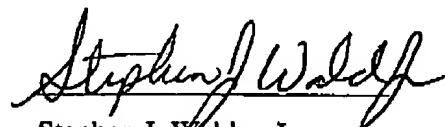
features are explicitly taught or suggested in the references. To the contrary, the Office Action makes irrelevant allegations regarding the Lee reference and then somehow concludes that the features of the claims are obvious even though not so much as a single feature of the claims is shown to be taught or suggested by Lee. Lee does not teach or suggest any of the features of claims 2-3, 10-11 and 18-19 and therefore, Applicants respectfully request withdrawal of the rejection of these claims under 35 U.S.C. § 103(a).

IV. Conclusion

It is respectfully urged that the subject application is patentable over Lee and is now in condition for allowance. The Examiner is invited to call the undersigned at the below-listed telephone number if in the opinion of the Examiner such a telephone conference would expedite or aid the prosecution and examination of this application.

Respectfully submitted,

DATE: March 29, 2004



Stephen J. Walder, Jr.
Reg. No. 41,534
Carstens, Yee & Cahoon, LLP
P.O. Box 802334
Dallas, TX 75380
(972) 367-2001
Attorney for Applicants